

(19)



Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 932 282 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
28.07.1999 Bulletin 1999/30

(51) Int. Cl.⁶: H04L 12/56

(21) Application number: 99300497.7

(22) Date of filing: 25.01.1999

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(72) Inventors:
• Chapman, Alan Stanley John
Kanata, Ontario K2K 1V5 (CA)
• Kung, Hsiang-Tsung
Lexington, MD 02173 (US)

(30) Priority: 27.01.1998 US 14110


(74) Representative: Cage, John David
Nortel Networks, IP Law Group,
London Road
Harlow, Essex CM17 9NA (GB)

(71) Applicant:
NORTEL NETWORKS CORPORATION
Montreal, Quebec H2Y 3Y4 (CA)

(54) TCP admission control

(57) Congestion at a network node can be aggravated by having too many TCP connections. A method of avoiding the bad effects of too many TCP connections is to limit the number of connections. Limiting the number of connections is achieved by an admission control (10) which delays or even discards the connec-

tion set-up packets. TCP traffic flows are monitored to generate packet loss characteristics and when a certain condition is met, a connection request queue is disabled.



EP 0 932 282 A2

Description

Field of the Invention

[0001] The invention relates generally to traffic congestion management of a data network. In particular, it is directed to a technique by which congestion in the data network is controlled by limiting new TCP connection setups based on packet loss characteristics of the data network.

Background of the Invention

[0002] The current data networks are handling not only enormous volume of traffic but more and more diversified multi media traffic, causing the data network to become congested more often. When congestion causes an excessive number of packets to be dropped, it can easily impact many traffic flows, and cause many timeouts. By guaranteeing a certain number of traffic flows a minimum bandwidth and treating the remainder as best effort, it is possible to avoid spreading high packet loss over so many flows and to reduce the number of aborted flows. Pending U. S. Patent Application Serial Nos. 08/772,256 filed on Dec. 23, 1996 and 08/818,612 filed on Mar. 14, 1997 by the present inventors describe dynamic traffic conditioning techniques which make use of this concept. The dynamic traffic conditioning techniques described therein allow the network to discover the nature of the service for each traffic flow, classify it dynamically, and exercise traffic conditioning by means of such techniques as admission control and scheduling when delivering the traffic downstream to support the service appropriately.

[0003] Congestion at a network node can be aggravated by having too many TCP connections. TCP will adjust to try to share bandwidth among all connections but when the available buffer space is insufficient, timeouts will occur and as the congestion increases there will be an exponentially growing number of packets resent. The effect of having too many connections is that much of the bandwidth in the upstream network is wasted carrying packets that will be discarded at the congested node because there is not enough buffer there.

[0004] A simple method of avoiding the bad effects of too many TCP connections is to limit the number of connections or to discard one or more packets from one or more existing connections. Limiting the number of connections is achieved by an admission control which delays or even discards the connection set-up packets. In the case of discarding packets, which packets and from which connection to discard packets are decided by preset algorithms or policies. By invoking this control to limit the number of connections, each packet is inspected to see if it is a connection set-up packet, e.g., TCP SYN packet. This control packet is used to initiate a TCP connection and no traffic can flow until it is

acknowledged by the other end of the proposed connection.

Summary of the Invention

[0005] It is an object of the invention to obviate disadvantages of the known methods.

[0006] It is a further object of the invention to provide a method of managing a data network for congestion.

[0007] It is a further object of the invention to provide a method of continuously monitoring the TCP traffic flows for congestion in a data network.

[0008] It is another object of the invention to provide a method of managing the data network by performing admission control for TCP traffic.

[0009] It is yet another object of the invention to provide a method of managing the data network by exercising the connection admission control for a new TCP connection request based on the packet loss characteristic.

[0010] Briefly stated, the invention resides in a packet data network for multimedia traffic having one or more nodes in which network one or more packets are discarded to control congestion. According to one aspect, a method of performing admission control to connection oriented traffic flows comprises steps of monitoring packets of all the traffic flows, deriving a packet loss characteristic of the traffic flows and disabling the serving of a new connection request when the packet loss characteristic matches a predefined pattern.

[0011] In another aspect, a method of performing admission control to TCP traffic flows comprises steps of storing all TCP connection setup packets in a connection request queue, monitoring packets of all active TCP traffic flows according to their port numbers and sequence numbers, and recording the count of either resent or discarded packets for any TCP traffic flows. The method further includes steps of building a history table containing the history of the sequence numbers, port numbers, and the count of either resent or discarded packets, computing a packet loss characteristic using the contents of the history table, and deciding enabling or disabling the connection request queue based on the packet loss characteristic with respect to a predefined pattern.

[0012] In a further aspect, the invention is directed to a TCP admission control apparatus for controlling congestion of a data network. The apparatus comprises a TCP output buffer for buffering and inspecting all the TCP packets of an incoming traffic flow, and a connection request queue for storing new connection requests. The apparatus further includes a history table for storing traffic information with respect to the TCP packets inspected above to derive a packet loss characteristic, and a queue controller for enabling or disabling the connection request queue upon detecting the matching of the packet loss characteristic with a predefined pattern.

Brief Description of Drawings

[0013]

Figure 1 is a schematic diagram of the admission control according to an embodiment of the invention.

Figures 2a and 2b are a flow chart for the case where TCP admission control is applied in a traffic link.

Figure 3 illustrates the relationship of admission control with the traffic conditioner.

Figures 4a and 4b are a flow chart for the case where TCP admission control is applied in a router.

Figures 5 and 6 show possible locations of admission control of the invention.

Detailed Description of the Preferred Embodiments of the Invention

[0014] Referring to Figure 1, the TCP admission control apparatus 10, according to one embodiment of the invention, includes a connection request queue 12. It is located at or near the output buffer 14 of a node of a data network. It should be noted that an admission control apparatus can be a separate device or can be made integral with or to reside in any node or link equipment. It should also be understood that TCP traffic flows as a whole can be processed by an apparatus or separate apparatus can be provided for each traffic flow or a group of traffic flows in one class. Every packet of an input stream is inspected and TCP packets are identified at the output buffer 14 using, for example, source and destination IP addresses, source and destination port numbers and protocol. All new connection requests are read at a connection reader 16 and are stored at the connection request queue 12. The connection request queue 12 is a FIFO. If admission control is not invoked then the new connection requests will be served immediately by enabling the connection request queue. If admission control is switched on then they will be delayed.

[0015] The admission control detects the packets that are being discarded and looks for multiple successive packets from the same flow or multiple instances of the same packet, the latter being the result of packet resends due to packet loss or discard. The admission control derives some pattern of packet discards by using a discard measure. For convenience, this measure is called packet loss characteristic in this specification. It is possible that other parameters can be used to indicate the state of congestion in a data network. If certain criteria are met or the packet loss characteristic matches a predefined pattern, admission control is

invoked and any new connection requests (connection set-up packets) will be delayed by disabling the connection request queue or packets belonging to one or more existing connections will be discarded until the problem clears. If a connection set-up packet is delayed too long (e.g., one second), it will be discarded from the queue.

[0016] When the packet loss characteristic shows that new connections can be accepted the servicing of the connection request queue is enabled. Waiting connection requests can be served immediately or can be released at a controlled pace according to a predefined algorithm.

[0017] The admission control apparatus therefore includes a small history table 18 and information about discarded packets is entered into it. When a packet is discarded, the flow identity (source and destination IP plus TCP socket number) is extracted and compared with current table entries. If the flow already has an entry then the history is updated. If the flow does not have an entry and there is room for a new entry, the new entry is made. If there is no room for a new entry the information is discarded.

[0018] The admission control can be performed on a traffic link or at a router.

[0019] In the case where the admission control is performed on the traffic link, the history table contains, for each active flow (or as many flows as can be handled), the following entries:

[0020] The first entry is a count of resent packets for that flow (Total Packet Resent).

[0021] The second entry is a count of how many times the currently recorded packet (that is the currently stored sequence number) has been resent (Same Packet Resent).

[0022] The third entry is the time that the most recent update was made for that flow. After some period of inactivity the flow is taken out of the table.

[0023] This information is used to look for patterns of discard that indicate congestion problems. It is assumed that if the sequence number on an arriving packet is lower than or equal to the stored value, then it must be a resend. The total number of resends as a fraction of the total number of packets is a measure of downstream congestion. In this embodiment, this measure is used as the packet loss characteristic.

[0024] Seeing the same packet resent multiple times will suggest that the connection is experiencing time-out or at least a very high loss rate. It is not usual for a packet to be discarded multiple times. Normally the TCP protocol will adjust its window to fit the available bandwidth and will only lose one packet before reducing that window. Although TCP relies on packet loss to constantly test for available bandwidth, a packet that is discarded once will almost certainly be forwarded when it is retransmitted. Multiple instances of the same packet will suggest that the TCP source is experiencing time-out.

[0025] There will be many variations on what informa-

tion is stored and what algorithm is used to assess whether new connections should be enabled.

[0026] It is not necessary to keep information on all flows since a sampled history is sufficient to detect problem conditions.

[0027] Entries in the history table are removed after a period of time. Also, whenever admission control is invoked, the history table is cleaned out and starts fresh to get a good picture of the new loss characteristic. The history table would be purged, in any case, at regular intervals to keep the history reflecting current loss characteristics. The interval would be configurable depending on line rates and expected number of flows, etc.

[0028] Figures 2a and 2b are a flow chart for the case where TCP admission control is applied in a traffic link rather than in a router.

[0029] As mentioned earlier, the applicant's pending applications describe traffic conditioners and Figure 3 shows one of such conditioners. In the Figure, a traffic conditioner 40 includes a plurality of queues 42, at least one for each class of TCP traffic. Every packet of an input stream is inspected and identified at 44 using, for example, IP addresses, ports, etc. A controller 46 characterises the flow (using rate, duration, etc.) and assigns it a class. The controller refers to a database 48 and uses output scheduling to allocate bandwidth among classes. It can implement an admission control policy of the present invention for a class before delivering an output stream toward downstream nodes or to peripherals. In this case it is necessary to work out whether a packet has been discarded, by looking for a second copy of it passing through the link.

[0030] In another embodiment, the admission control is performed in the router where the discarded packets can be inspected directly as the discard decision is made at the buffer of the router.

[0031] In this case the history table contains, for each active flow (or as many flows as can be handled), the following entries:

[0032] The first entry is a count of discarded packets for that flow (Total Packet Discarded).

[0033] The second entry is a count of how many times the currently recorded packet (that is the currently stored sequence number) has been discarded (Same Packet Discarded).

[0034] The third entry is the time that the most recent update was made for that flow. After some period of inactivity the flow is taken out of the table.

[0035] This information is used to look for patterns of discard that indicate congestion problems. The total number of discards as a fraction of the total number of packets is a measure of buffer congestion.

[0036] Seeing the same packet resent multiple times will suggest that the connection is experiencing time-out or at least a very high loss rate.

[0037] There will be many variations on what information is stored and what algorithm is used to assess whether new connections should be enabled.

[0038] In another embodiment, if the admission control is performed at the router, packets from one or more existing connections can be discarded to control congestion at its buffer. The discarding action can be taken together with action of limiting the set-up of new connections, latter having been described above.

[0039] Figures 4a and 4b are a flow chart for the case where TCP admission control is applied in a router rather than in a traffic link.

[0040] Like the traffic conditioning of the pending applications, the admission control can take place at various places in the data network and can be biased toward certain kinds of TCP traffic. For example, as gateways are often a bottleneck and bulk flows can decrease response times for interactive users, an admission control can be located at a place shown in Figure 5 which will alleviate this problem. In Figure 6, traffic conditioners are located at a plurality of IP switches which form a data network 60

[0041] In summary, congestion at a network node can be aggravated by having too many TCP connections. A simple method of avoiding the bad effects of too many TCP connections is to limit the number of connections. Limiting the number of connections is achieved by an admission control which delays or even discards the connection set-up packets. TCP traffic flows are monitored to generate packet loss characteristics and when a certain condition is met, a connection request queue is disabled.

Claims

1. A method of performing admission control to connection oriented traffic flows in a packet data network for multimedia traffic having one or more nodes in which network one or more packets are discarded to control congestion, the method comprising the steps of;

monitoring packets of all the traffic flows;

deriving a packet loss characteristic of the traffic flows; and

disabling the serving of a new connection request when the packet loss characteristic matches a predefined pattern.

2. The method of performing admission control to traffic flows according to claim 1 wherein the connection oriented traffic flows are TCP traffic flows and the step of deriving a packet loss characteristic comprises further steps of;

monitoring discarded packets for the TCP traffic flows;

generating a history table containing history of

the discarded packets for active TCP traffic flows; and

analysing the history table to derive the packet loss characteristic.

3. The method of performing admission control, according to claim 2 wherein the step of generating the history table comprising steps of:

entering a count of discarded packets for an active TCP flow, and

entering a count of how many times the currently recorded packet has been discarded.

4. The method of performing admission control to connection oriented traffic flows according to claim 1 wherein the connection-oriented traffic flows are TCP traffic flows and the step of deriving a packet loss characteristic comprises further steps of;

monitoring resent packets for TCP traffic flows;

generating a history table containing history of the resent packets for active TCP traffic flows; and

analysing the history table to derive the packet loss characteristic.

5. The method of performing admission control, according to claim 4 wherein the step of generating the history table comprises steps of:

entering a count of resent packets for an active TCP traffic flow, and

entering a count of how many times the currently recorded packet has been resent.

6. The method of performing admission control, according to claim 3 or 5 further comprising steps of:

storing new TCP connection requests in a connection request queue; and

clearing all the entries of the history table whenever the connection request queue is re-enabled.

7. The method of performing admission control, according to claim 3 or 5 further comprising steps of:

purging all the entries of the history table periodically from time to time or after a certain pre-

set period of time.

8. The method of performing admission control, according to claim 3 or 4, comprising a further step of enabling the serving of a plurality of new connection requests at a controlled pace.

9. A method of performing admission control to TCP traffic flows in a packet data network for multimedia traffic having one or more nodes in which network one or more packets are discarded to control congestion; the method comprising the steps of;

storing all TCP connection setup packets in a connection request queue;

monitoring packets of all active TCP traffic flows according to their port numbers and sequence numbers;

recording the count of either resent or discarded packets for any TCP traffic flows;

building a history table containing the history of the sequence numbers, port numbers, and the count of either resent or discarded packets;

computing a packet loss characteristic using the contents of the history table; and

deciding enabling or disabling the connection request queue based on the packet loss characteristic with respect to a predefined pattern.

10. The method of performing admission control to TCP traffic flows according to claim 9 wherein the step of computing a packet loss characteristic comprises step of:

deriving the total number of either resends or discards as a fraction of the total number of TCP packets of the TCP traffic flow.

11. The method of performing admission control to TCP traffic flows according to claim 10, comprising a further step of:

deciding to disable the connection request queue when the fraction reaches a preset threshold.

12. The method of performing admission control to TCP traffic flows according to claim 9, comprising a further step of:

enabling the connection request queue at a controlled pace.

13. A TCP admission control apparatus for controlling congestion of a data network, comprising:

a TCP output buffer for buffering and inspecting all the TCP packets of an incoming traffic flow; 5

a connection request queue for storing new connection requests;

a history table for storing traffic information with respect to the TCP packets inspected above to derive a packet loss characteristic; and 10

a queue controller for enabling or disabling the connection request queue upon detecting the matching of the packet loss characteristic with a predefined pattern. 15

14. The TCP admission control apparatus according to claim 13 wherein the history table contains entries of a count of either resent or discarded packets for the traffic flow and the total number of TCP packet of the TCP traffic flow. 20

25

30




35

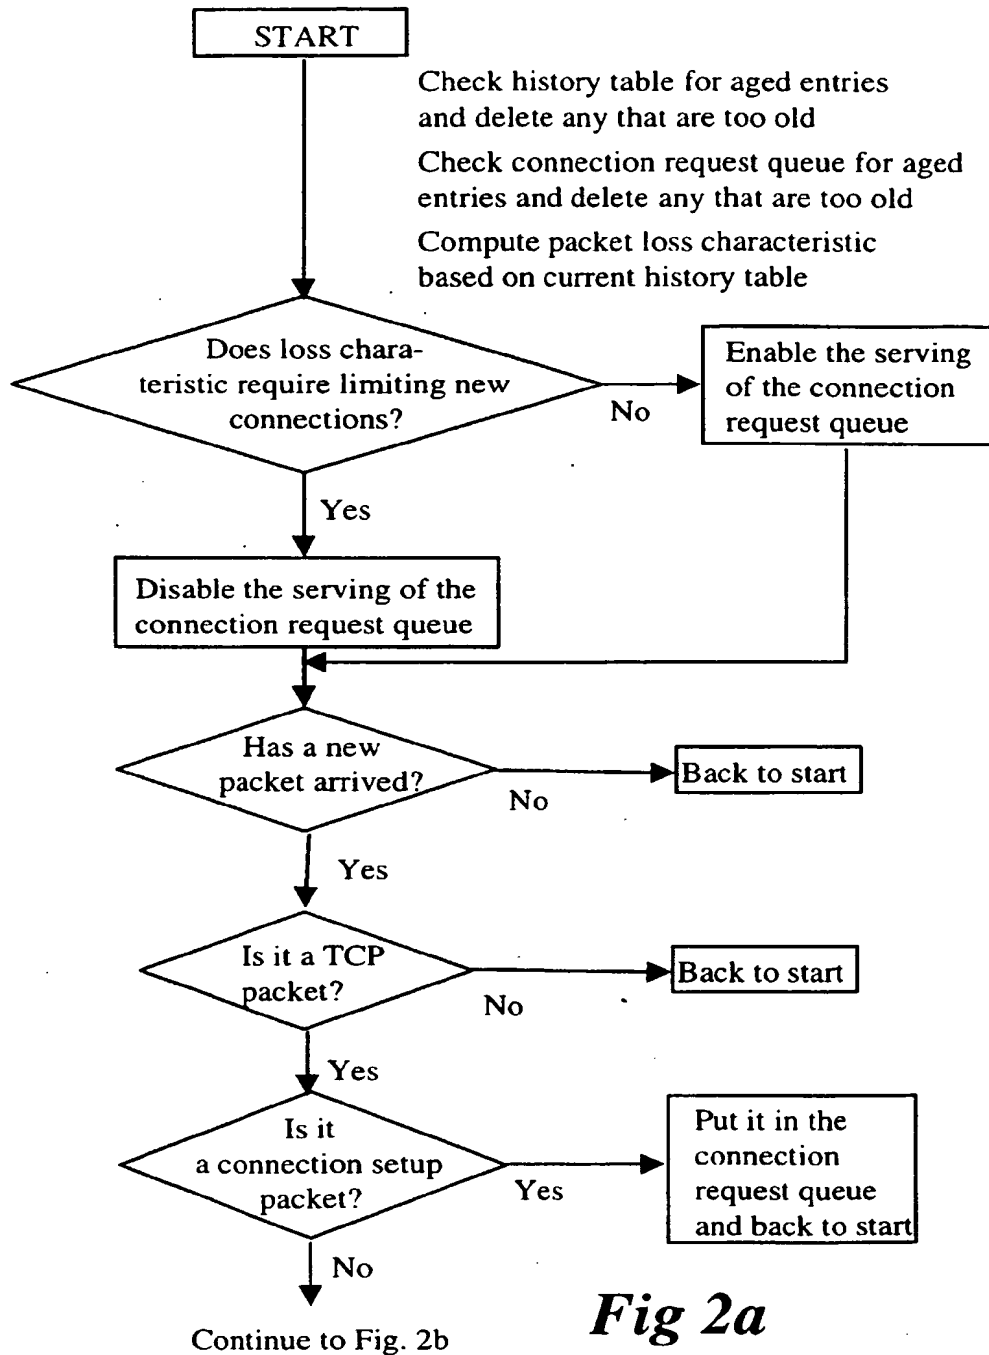
40

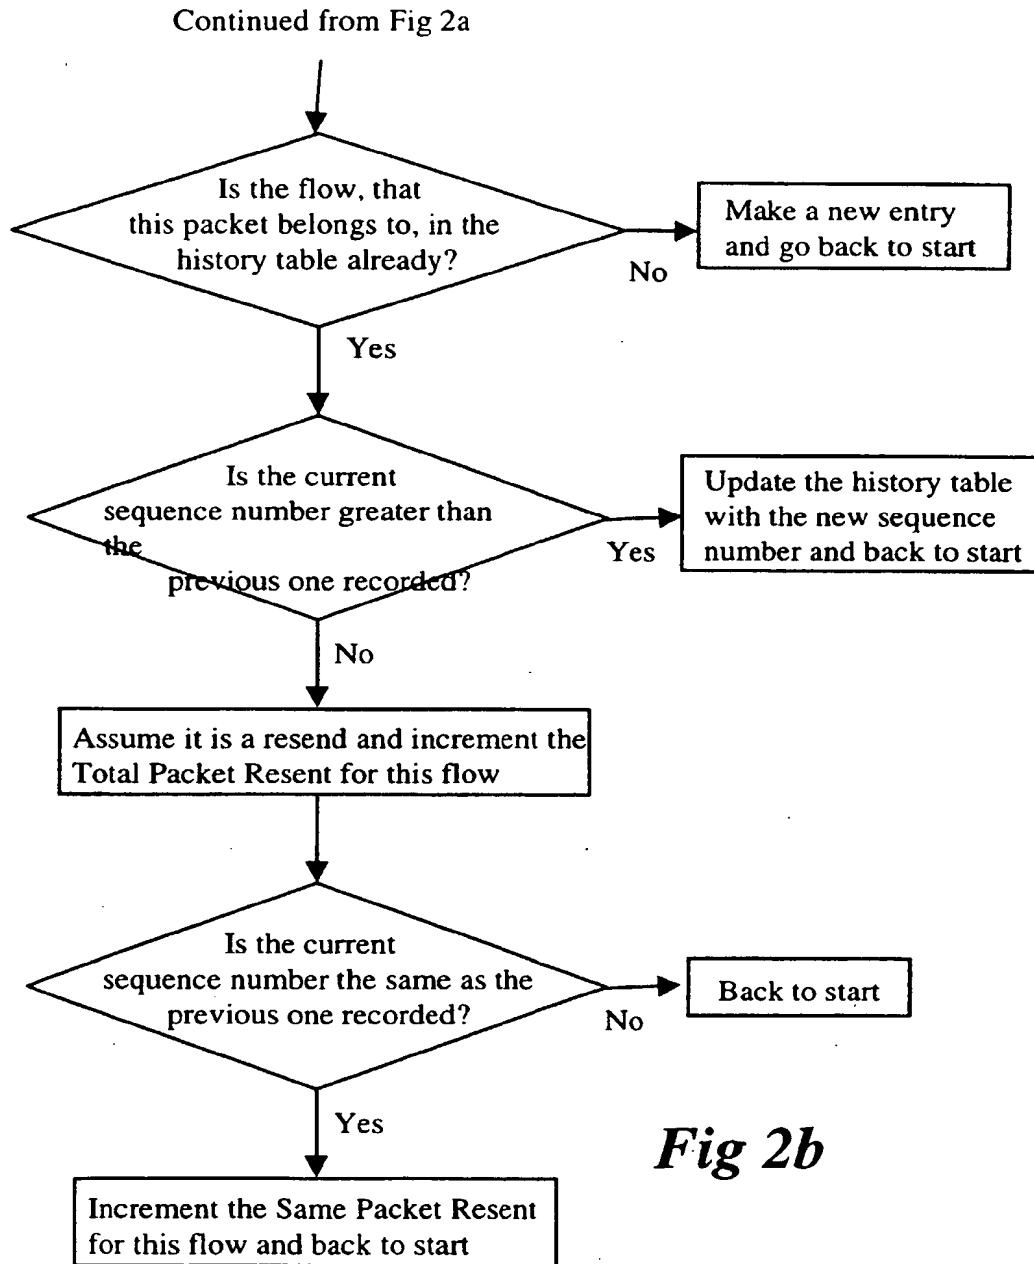
45

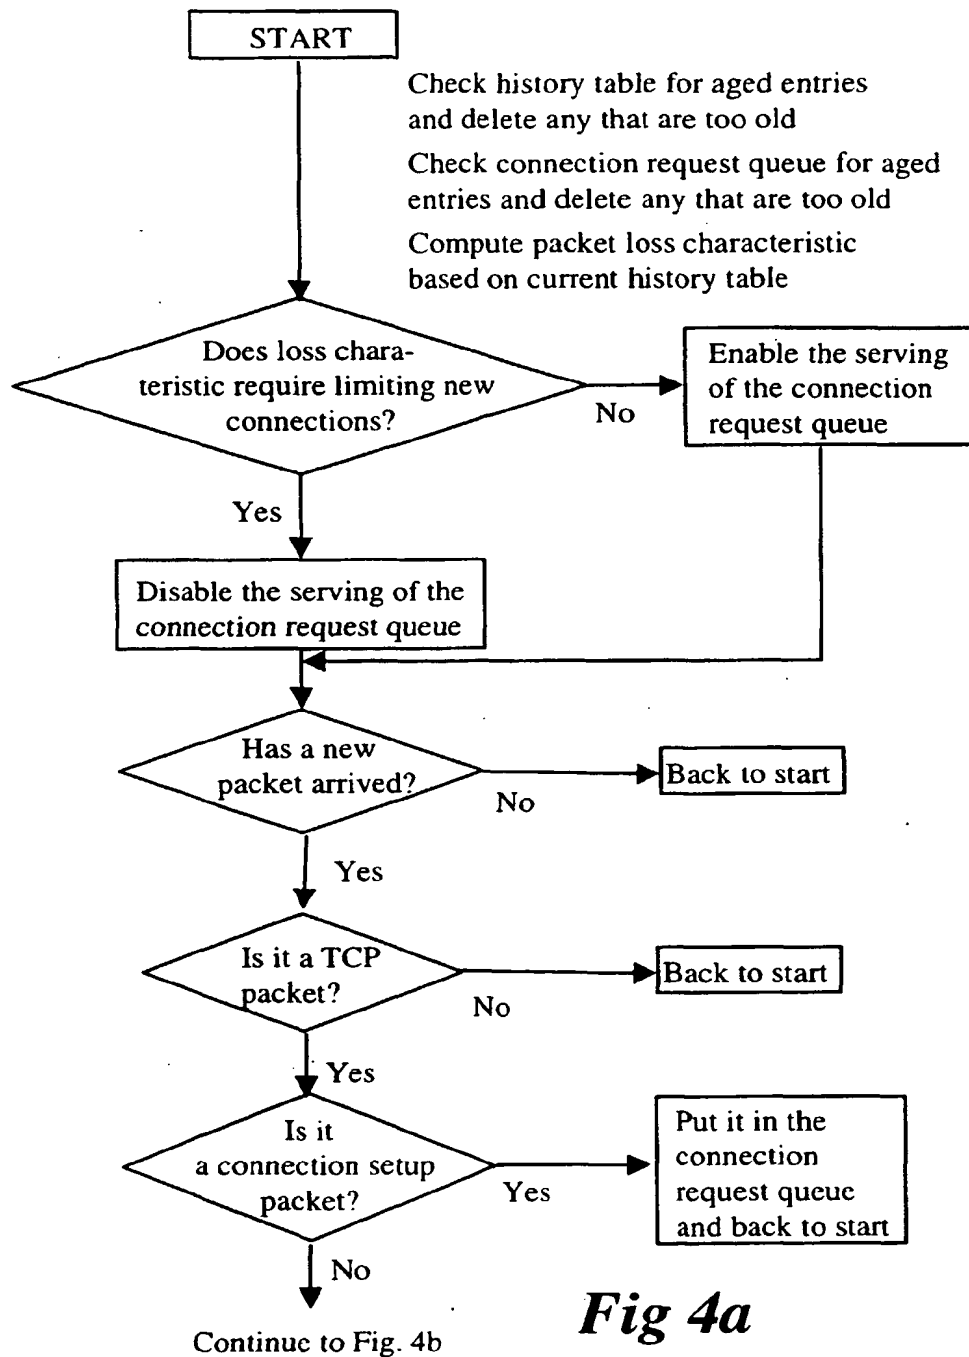
50

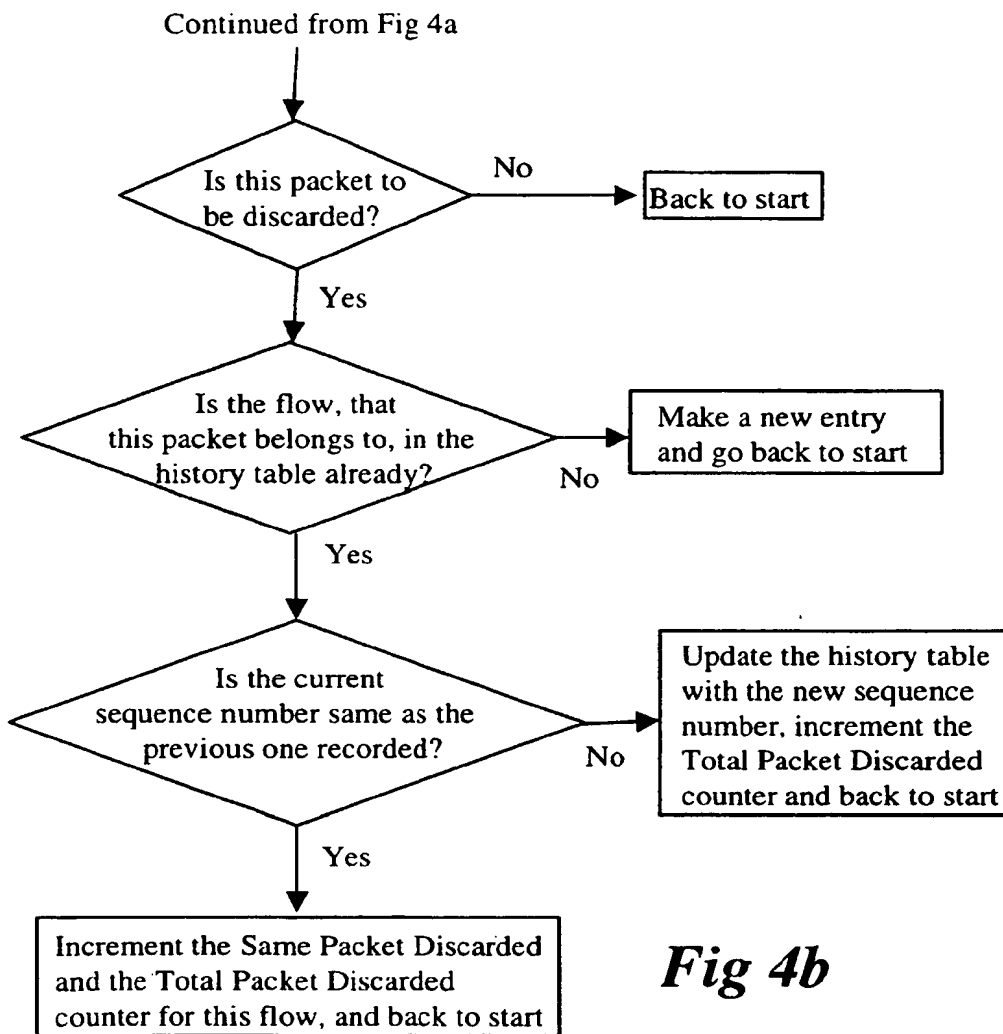
55





**Fig 2a**



**Fig 4a**







(19)

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11)

EP 0 932 282 A3

(12)

EUROPEAN PATENT APPLICATION

(88) Date of publication A3:
30.08.2000 Bulletin 2000/35

(51) Int. Cl.⁷: H04L 12/56

(43) Date of publication A2:
28.07.1999 Bulletin 1999/30

(21) Application number: 99300497.7

(22) Date of filing: 25.01.1999

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SE
Designated Extension States:
AL LT LV MK RO SI

(30) Priority: 27.01.1998 US 14110

(71) Applicant:
NORTEL NETWORKS CORPORATION
Montreal, Quebec H2Y 3Y4 (CA)


(72) Inventors:
• Chapman, Alan Stanley John
Kanata, Ontario K2K 1V5 (CA)
• Kung, Hsiang-Tsung
Lexington, MD 02173 (US)

(74) Representative: Cage, John David
Nortel Networks, IP Law Group,
London Road
Harlow, Essex CM17 9NA (GB)

(54) TCP admission control

(57) Congestion at a network node can be aggravated by having too many TCP connections. A method of avoiding the bad effects of too many TCP connections is to limit the number of connections. Limiting the number of connections is achieved by an admission control (10) which delays or even discards the connec-

tion set-up packets. TCP traffic flows are monitored to generate packet loss characteristics and when a certain condition is met, a connection request queue is disabled.



EP 0 932 282 A3



European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 99 30 0497

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
X	EP 0 473 188 A (TOKYO SHIBAURA ELECTRIC CO) 4 March 1992 (1992-03-04) * column 1, line 35 - line 50 * * figure 6 *	1	H04L12/56
A	---	2-14	
L	WO 99 66676 A (TURANYI ZOLTAN ;VERES ANDRAS (HU); ERICSSON TELEFON AB L M (SE)) 23 December 1999 (1999-12-23) * figures 2,4,5 * -----	1-14	
			TECHNICAL FIELDS SEARCHED (Int.Cl.6)
			H04L H04Q
The present search report has been drawn up for all claims			
Place of search BERLIN		Date of completion of the search 11 July 2000	Examiner Siebel, C
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons & : member of the same patent family, corresponding document</p>			

EPO FORM 1500 03.82 (P04C01)

**ANNEX TO THE EUROPEAN SEARCH REPORT
ON EUROPEAN PATENT APPLICATION NO.**

EP 99 30 0497

This annex lists the patent family members relating to the patent documents cited in the above-mentioned European search report.
The members are as contained in the European Patent Office EDP file on
The European Patent Office is in no way liable for these particulars which are merely given for the purpose of information.

11-07-2000

Patent document cited in search report		Publication date	Patent family member(s)	Publication date
EP 0473188	A	04-03-1992	JP 4115643 A	16-04-1992
			CA 2050436 A	01-03-1992
			DE 69114789 D	04-01-1996
			DE 69114789 T	09-05-1996
			US 5267232 A	30-11-1993

W0 9966676	A	23-12-1999	GB 2338372 A	15-12-1999
			AU 4364699 A	05-01-2000

EPO FORM P0499

For more details about this annex : see Official Journal of the European Patent Office, No. 12/82